

Density Ratio Hidden Markov Models

John A. Quinn

Department of Computer Science
Makerere University
PO Box 7062, Kampala, Uganda
jquinn@cit.ac.ug

Masashi Sugiyama

Department of Computer Science
Tokyo Institute of Technology
Tokyo 152-8552, Japan
sugi@cs.titech.ac.jp

Abstract

Hidden Markov models and their variants are the predominant sequential classification method in such domains as speech recognition, bioinformatics and natural language processing. Being generative rather than discriminative models, however, their classification performance is a drawback. In this paper we apply ideas from the field of density ratio estimation to bypass the difficult step of learning likelihood functions in HMMs. By reformulating inference and model fitting in terms of density ratios and applying a fast kernel-based estimation method, we show that it is possible to obtain a striking increase in discriminative performance while retaining the probabilistic qualities of the HMM. We demonstrate experimentally that this formulation makes more efficient use of training data than alternative approaches.

1 Introduction

Inference of a sequence of estimated classes from a sequence of noisy observations is fundamental in many applications. The hidden Markov model (HMM) and its variants are the usual methods employed to do this, and have been used with conspicuous success in such domains as speech recognition, bioinformatics and natural language processing. As well as being computationally efficient, they are a popular choice due to their intuitive probabilistic interpretation. However, they have drawbacks in terms of classification accuracy, being primarily generative rather than discriminative models.

One established approach to improve classification performance in HMMs models has been to adapt the model to discriminative forms which apply information theoretic principles in training [6, 22]. This improves classification performance, though the requirement in general to limit such models to parametric forms means they still do not have the discriminative power of kernel-based and max-margin methods. Another idea therefore is to create models combining the structure of the HMM and the classification approach of Support Vector Machines [1, 19, 20]. This considerably improves classification performance, though at the expense of losing the intuitive probabilistic interpretation of the HMM.

In this paper, we propose a different idea which combines the advantages of both the above approaches, using concepts from the field of direct density ratio estimation [16, 7, 2]. Our key observation is that rather than trying to quantify a set of *likelihoods* (i.e. how likely any possible observation is given some class, relative to the likelihood of making other observations given that same class), it is in fact only necessary to know about *likelihood ratios* (i.e. how likely any possible observation is given some class, relative to the likelihood of making the same observation given a different class). Because the forward-backward inference algorithm computes such ratios anyway, we therefore follow the principle that we should avoid solving a more general learning problem than is strictly necessary [21].

We therefore reformulate the forward-backward algorithm for HMM inference in terms of density ratios so that the intermediate step of likelihood function estimation can be dispensed with. We demonstrate how efficient and highly discriminative nonparametric inference can be carried out in this framework using a kernel-based density ratio estimation procedure [4]. Because density ratios are a natural parameterization of

the forward-backward algorithm, the resulting inference procedure is also more numerically stable than the conventional version. We demonstrate these ideas using synthetic data and on a physiological monitoring problem. In the case of physiological monitoring, we also demonstrate an application to sequential anomaly detection.

The structure of the rest of the paper is as follows. Section 2 reviews the conventional forward-backward algorithm for hidden Markov model inference. Section 3 recasts this procedure in terms of ratios of probability densities, and Section 4 describes how nonparametric estimates of those density ratios can be calculated directly from training data in both supervised and unsupervised settings, and how estimates of several pairwise likelihood ratio functions can be obtained from a concise set of parameters. Experimental results in Section 5 show the performance of the density ratio HMM methodology on synthetic and real-world physiological monitoring data, showing striking improvements compared to conventional parametric and nonparametric sequential inference approaches.

A Matlab implementation and demo is available at <http://cit.mak.ac.ug/staff/jquinn/software/densityratioHMM.html>.

2 Forward-backward inference

Consider the estimation of a latent sequence $x_{1:T} = \{x_t \in \mathcal{S} | t = 1, \dots, T\}$ from an observation sequence $\mathbf{y}_{1:T} = \{\mathbf{y}_t \in \mathbb{R}^{d_y} | t = 1, \dots, T\}$. The variable x_t is assumed to have first order Markovian dynamics, such that $p(x_t | x_{1:t-1}) = p(x_t | x_{t-1})$, the values it can take on are a discrete set of classes or ‘states’ $\mathcal{S} = \{1, \dots, S\}$, and each observation is independently drawn from a fixed emission distribution $p(\mathbf{y}_t | x_t)$. Given a sequence $\mathbf{y}_{1:T}$ and knowledge of both $p(x_t | x_{t-1})$ and $p(\mathbf{y}_t | x_t)$ and the initial state distribution $p(x_1)$, the forward-backward algorithm can be used to estimate the probability of each state at each time frame, $p(x_t = i | \mathbf{y}_{1:T})$. We use the following shorthand in describing the algorithm: \mathbf{A} denotes a matrix of state transition probabilities, such that $\mathbf{A}_{ij} = p(x_t = j | x_{t-1} = i)$ and π is a vector of initial state probabilities such that $\pi_i = p(x_1 = i)$.

The first stage of the algorithm involves recursively calculating forward messages $\alpha_t(i)$ for each of the states $i \in \mathcal{S}$:

$$\tilde{\alpha}_i(1) = \pi_i p(\mathbf{y}_1 | x_1 = i), \quad (1)$$

$$\tilde{\alpha}_i(t) = \left[\sum_{j \in \mathcal{S}} \alpha_j(t-1) \mathbf{A}_{ji} \right] p(\mathbf{y}_t | x_t = i), \quad (2)$$

$$t = 2, \dots, T.$$

After each time step, it is conventional (but not necessary) to normalize the forward messages,

$$\alpha_i(t) = c_t \tilde{\alpha}_i(t) \quad \text{s.t.} \quad \sum_{i \in \mathcal{S}} \alpha_i(t) = 1. \quad (3)$$

Without this normalization the algorithm is numerically unstable, as the values of $\tilde{\alpha}_i$ otherwise become very small over repeated iterations. The step is optional however because normalization occurs anyway later in the procedure (in Eq. (7)). The forward messages can be interpreted as the posterior probability of each state given observations up to that time frame, the process of calculating this being known as *filtering*. Note that the normalization means that the absolute values of $p(\mathbf{y}_t | x_t = i)$ are not directly significant for inference; we are ultimately interested only in the relative magnitudes.

To carry out *smoothing* (calculation of $p(x_t = i | \mathbf{y}_{1:T})$) the backwards messages $\beta_t(i)$ must first be similarly calculated:

$$\tilde{\beta}_i(T) = 1, \quad (4)$$

$$\tilde{\beta}_i(t) = \sum_{j \in \mathcal{S}} \mathbf{A}_{ij} p(\mathbf{y}_{t+1} | x_{t+1} = j) \beta_j(t+1), \quad (5)$$

$$t = T-1, \dots, 1,$$

also with an optional normalization step carried out after each iteration,

$$\beta_i(t) = c'_t \tilde{\beta}_i(t) \quad \text{s.t.} \quad \sum_{i \in \mathcal{S}} \beta_i(t) = 1. \quad (6)$$

The two types of messages are then combined to give the final result

$$\gamma_i(t) = p(x_t=i|\mathbf{y}_{1:T}) = \frac{\alpha_i(t)\beta_i(t)}{\sum_{j \in \mathcal{S}} \alpha_j(t)\beta_j(t)}. \quad (7)$$

The forward-backward algorithm therefore requires an explicit likelihood model for every dynamical regime, is numerically unstable (requiring message scaling or transforms in and out of log space to prevent underflow), and loses information during normalization steps. We next describe how a density ratio formulation overcomes these problems, by expressing inference in terms of parameters which are more natural to the problem.

3 Inference with density ratios

In this section we rearrange the above inference equations in terms of pairwise probability density ratios. Ratios of probabilities make it particularly convenient to express Bayesian updates, as the ratio of two posteriors is equal to simply the ratio of the priors multiplied by the ratio of the likelihoods. In order to express the forward-backward equations in this way, corresponding to the three types of values $\alpha_i, \beta_i, \gamma_i$ we define three types of ratios $\vec{r}_{i,j}, \overleftarrow{r}_{i,j}, r_{i,j}$ as

$$\begin{aligned} \vec{r}_{i,j}(t) &= \frac{\alpha_i(t)}{\alpha_j(t)} = \frac{p(x_t=i|\mathbf{y}_{1:t})}{p(x_t=j|\mathbf{y}_{1:t})}, \\ \overleftarrow{r}_{i,j}(t) &= \frac{\beta_i(t)}{\beta_j(t)} = \frac{p(x_t=i|\mathbf{y}_{t+1:T})}{p(x_t=j|\mathbf{y}_{t+1:T})}, \\ r_{i,j}(t) &= \frac{\gamma_i(t)}{\gamma_j(t)} = \frac{p(x_t=i|\mathbf{y}_{1:T})}{p(x_t=j|\mathbf{y}_{1:T})}. \end{aligned}$$

We also treat likelihoods in terms of density ratios, and define

$$w_{i,j}(\mathbf{y}_t) = \frac{p(\mathbf{y}_t|x_t=i)}{p(\mathbf{y}_t|x_t=j)}.$$

Using Eqs. (1-2) we can derive the following expressions for density ratio versions of the forward equations:

$$\vec{r}_{i,j}(1) = \frac{\pi_i}{\pi_j} w_{i,j}(\mathbf{y}_1), \quad (8)$$

$$\begin{aligned} \vec{r}_{i,j}(t) &= \frac{[\sum_{k \in \mathcal{S}} \alpha_k(t-1) \mathbf{A}_{ik}] p(\mathbf{y}_t|x_t=i)}{[\sum_{k' \in \mathcal{S}} \alpha_{k'}(t-1) \mathbf{A}_{jk'}] p(\mathbf{y}_t|x_t=j)} \\ &= \sum_{k \in \mathcal{S}} \left[\sum_{k' \in \mathcal{S}} \frac{\mathbf{A}_{jk'}}{\mathbf{A}_{ik}} \vec{r}_{k',k}(t-1) \right]^{-1} w_{i,j}(\mathbf{y}_t) \\ & \quad t = 2, \dots, T. \end{aligned} \quad (9)$$

Although these expressions are written in terms of all pairwise ratios, implying that S^2 terms have to be calculated at each time frame, in fact these ratios can be estimated with a concise set of parameters as we demonstrate in Section 4.

For backwards messages we use Eqs. (4-5) in a similar way to obtain

$$\overleftarrow{r}_{i,j}(T) = 1 \quad (10)$$

$$\begin{aligned} \overleftarrow{r}_{i,j}(t) &= \frac{\sum_{k \in \mathcal{S}} \mathbf{A}_{ki} p(\mathbf{y}_{t+1} | x_{t+1}=k) \beta_k(t+1)}{\sum_{k' \in \mathcal{S}} \mathbf{A}_{k'j} p(\mathbf{y}_{t+1} | x_{t+1}=k') \beta_{k'}(t+1)} \\ &= \sum_{k \in \mathcal{S}} \left[\sum_{k' \in \mathcal{S}} \frac{\mathbf{A}_{k'j}}{\mathbf{A}_{ki}} \overleftarrow{r}_{k',k}(t+1) w_{k',k}(\mathbf{y}_{t+1}) \right]^{-1}. \\ t &= T-1, \dots, 1. \end{aligned} \quad (11)$$

The messages are finally combined simply with

$$r_{i,j}(t) = \overrightarrow{r}_{i,j}(t) \overleftarrow{r}_{i,j}(t).$$

The filtering and smoothing probabilities, if needed, can be simply calculated from the ratios with $\alpha_i(t) = \frac{\overrightarrow{r}_{i,S}(t)}{\sum_{j \in \mathcal{S}} \overrightarrow{r}_{j,S}(t)}$ and $\gamma_i(t) = \frac{r_{i,S}(t)}{\sum_{j \in \mathcal{S}} r_{j,S}(t)}$. This completes the density ratio forward-backward algorithm.

To provide some intuition about the ratio-based procedure we look in some more detail at a two-state HMM, $\mathcal{S} = \{1, 2\}$. The forward equation in this simplified case reduces to:

$$\overrightarrow{r}_{1,2}(t) = \frac{\mathbf{A}_{11} \overrightarrow{r}_{1,2}(t-1) + \mathbf{A}_{12} w_{1,2}(\mathbf{y}_t)}{\mathbf{A}_{21} \overrightarrow{r}_{1,2}(t-1) + \mathbf{A}_{22} w_{1,2}(\mathbf{y}_t)}, \quad (12)$$

and the backward equation to

$$\overleftarrow{r}_{1,2}(t) = \frac{\mathbf{A}_{11} \overleftarrow{r}_{1,2}(t+1) w_{1,2}(\mathbf{y}_{t+1}) + \mathbf{A}_{21}}{\mathbf{A}_{21} \overleftarrow{r}_{1,2}(t+1) w_{1,2}(\mathbf{y}_{t+1}) + \mathbf{A}_{22}}. \quad (13)$$

The first term of (12) specifies the evolution from $\frac{p(x_{t-1}=1|\mathbf{y}_{1:t-1})}{p(x_{t-1}=2|\mathbf{y}_{1:t-1})}$ to $\frac{p(x_t=1|\mathbf{y}_{1:t-1})}{p(x_t=2|\mathbf{y}_{1:t-1})}$ in terms of the transition probabilities in \mathbf{A} . Hence this is simply the ratio form of marginalizing out the transition probabilities. Figure 1 shows this evolution of the state probability ratio given transition matrices of the form $\mathbf{A} = \begin{bmatrix} \alpha & 1-\alpha \\ 1-\alpha & \alpha \end{bmatrix}$

for different values of α . If $\alpha = 0.5$, then any past information is lost; $\frac{p(x_t=1|\mathbf{y}_{1:t-1})}{p(x_t=2|\mathbf{y}_{1:t-1})}$ is always equal to 1, hence both states are equally likely *a priori* at time t . If $\alpha = 1$, there can be no state transitions. Inference is equivalent to accumulating evidence in favor of either state 1 or 2 globally, and Eq. (12) reduces to $\overrightarrow{r}_{1,2}(t) = \frac{\pi_1}{\pi_2} \prod_{i=1}^t w_{1,2}(\mathbf{y}_i)$. The ratio form of the Bayesian update given the observation \mathbf{y}_t is simply multiplication by $w_{1,2}(\mathbf{y}_t)$.

The ratio formulation of the forward-backward algorithm is numerically stable, and Figure 1 provides an intuition of this; extreme ratio values tend to be mapped back to within a few orders of magnitude of unity after considering transition probabilities. Scaling of values to prevent underflow is therefore not necessary with this method as it is in the conventional forward-backward algorithm.

Hence, given a way to estimate $w_{i,j}(\cdot)$ from training data, it is possible to carry out inference in the HMM without ever needing to calculate the individual likelihoods. These steps are mathematically equivalent to the standard forward-backward algorithms, but do not require any normalization step. If the observation distribution is known exactly then the ratio formulation is equivalent and has no advantage. However, when the observation distribution needs to be approximated in some way – which is almost always the case with complex real-world data – the parameterization in terms of $w_{i,j}(\cdot)$ is more natural to the sequential inference problem than attempting to approximate each of the $p(\mathbf{y}_t | x_t=i)$ distributions directly. We discuss ways in which the estimation of these ratios and other parameters can be carried out in the next section.

4 Parameter estimation

The parameters to be learned in the density ratio HMM model are the likelihood ratio functions $w_{ij}(\mathbf{y})$, the transition matrix \mathbf{A} and initial class probabilities π . In this section we first suggest effective methods




Figure 1: Forward probability ratio evolution function given different transition matrices, for a 2-state HMM. See text for details.

for estimating $w_{ij}(\mathbf{y})$ which bypass the difficult step of estimating the densities $p(\mathbf{y}|x=i)$ and $p(\mathbf{y}|x=j)$ individually. In Section 4.1 we introduce an efficient method for carrying this out when there are two classes to be modeled, and give a related method for several classes in Section 4.2 which does not need every pairwise ratio to be estimated individually. The estimation of \mathbf{A} and π is then discussed in Section 4.3.

4.1 Two-class density ratio estimation

We begin by discussing how to estimate a likelihood ratio from data in the case that there are only two states i and j in the model, for example as would be needed to calculate (12) and (13). There is just one ratio to learn, since $w_{ij} = w_{ji}^{-1}$. Several techniques have been developed for direct density ratio estimation [10, 2, 3, 7, 18, 17]. We use a least squares approach here [4] which yields a consistent estimator with very good computational efficiency. The estimator is of the following form:

$$\hat{w}_{ij}(\mathbf{y}) = \boldsymbol{\theta}_{ij}^\top \boldsymbol{\phi}(\mathbf{y}),$$

where

$$\boldsymbol{\theta}_{ij} = (\theta_{ij,1}, \dots, \theta_{ij,B})^\top \in \mathbb{R}^B$$

for some number of parameters B , and

$$\boldsymbol{\phi}(\mathbf{y}) = (K(\mathbf{y}, \mathbf{y}_1), \dots, K(\mathbf{y}, \mathbf{y}_B))^\top \in \mathbb{R}^B \quad (14)$$

is a vector of kernel basis functions. We can set $B = N$ to have a kernel basis function at every training point, or for $B < N$ use some random subset of the training points. In this work we use the squared exponential kernel $K(\mathbf{y}, \mathbf{y}') = \exp\left(\frac{-\|\mathbf{y} - \mathbf{y}'\|^2}{2\sigma^2}\right)$.

This model can be fitted using a squared loss objective function,

$$J_{ij}(\boldsymbol{\theta}_{ij}) = \frac{1}{2} \int \left(\hat{w}_{ij}(\mathbf{y}) - \frac{p(\mathbf{y}|x=i)}{p(\mathbf{y}|x=j)} \right)^2 p(\mathbf{y}|x=j) d\mathbf{y}.$$

Expanding the squared term we obtain

$$\begin{aligned} J_{ij}(\boldsymbol{\theta}_{ij}) &= \frac{1}{2} \int \hat{w}_{ij}(\mathbf{y})^2 p(\mathbf{y}|x=j) d\mathbf{y} \\ &\quad - \int \hat{w}_{ij}(\mathbf{y}) p(\mathbf{y}|x=i) d\mathbf{y} + C, \end{aligned}$$

where C is a constant term that does not depend on any of the θ_{ij} values. Empirically, we can approximate the expectations by sample averages. Ignoring the constant C , factor $1/N$ and including an ℓ_2 -regularizer, we have the following training criterion:

$$\hat{J}_{ij}(\theta_{ij}) = \frac{1}{2} \theta_{ij}^\top \Phi^\top \mathbf{M}_j \Phi \theta_{ij} - \theta_{ij}^\top \Phi \mathbf{m}_i + \frac{\rho}{2} \|\theta_{ij}\|^2,$$

where $\Phi = (\phi(\mathbf{y}_1), \dots, \phi(\mathbf{y}_N))^\top$, \mathbf{m}_i is a column vector indicating membership of class i such that the j th element is one if $x_j = i$ and zero otherwise, and \mathbf{M}_j is a square matrix with \mathbf{m}_j along the diagonal and other entries set to zero. $\hat{J}_i(\theta_i)$ is minimized by

$$\hat{\theta}_{ij} = \left(\Phi^\top \mathbf{M}_j \Phi + \rho \mathbf{I}_B \right)^{-1} \Phi^\top \mathbf{m}_i. \quad (15)$$

We select ρ and σ with cross validation. Because of the nature of the estimator, it is sometimes possible to obtain negative values for $\hat{w}_{ij}(\mathbf{y})$. We simply round up negative estimates to zero in such cases, which does not affect the consistency of the estimator [5]. This least-squares approach is very fast to compute in practice, finding a global optimum in a single step with no iterative parameter search required.

4.2 Extension to multiple classes

In problems where $S > 2$, it would be possible to simply learn several pairwise likelihood ratio functions. Because $w_{ij} = w_{ji}^{-1}$ and $w_{ii} = 1$, this would entail learning $\frac{S^2 - S}{2}$ ratio functions. Fewer still need to be estimated if relationships of the form $w_{ij} = w_{ik}/w_{jk}$ are used, though this ‘ratio of ratios’ may become unstable when the denominator w_{jk} is close to zero.

For multiple-class ratio estimation another approach is to directly estimate conditional probabilities of class x_t given i.i.d. samples $\{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)}\}$. Using the fact that $p(x=i|\mathbf{y}) = \frac{p(\mathbf{y}|x=i)p(x=i)}{\sum_{j \in S} p(\mathbf{y}|x=j)p(x=j)}$, likelihood ratios in our problem can be estimated as

$$\hat{w}_{i,j}(\mathbf{y}) = \frac{p(\mathbf{y}|x=i)}{p(\mathbf{y}|x=j)} = \frac{p(x=j)}{p(x=i)} \frac{p(x=i|\mathbf{y})}{p(x=j|\mathbf{y})}. \quad (16)$$

This is similar in principle to the scheme of using classifiers to calculate likelihoods in standard HMMs discussed in (author?) [9].

A standard approach to calculating the $p(x_t=i|\mathbf{y})$ terms is kernel logistic regression. This can be computationally demanding for large datasets though, so we use a least squares procedure [15] similar in form to the two-class estimator we introduced in the previous section, and which is known to give comparable accuracy to kernel logistic regression but requiring far less training time.

We construct functions $q(x=i|\mathbf{y}, \theta_i)$ to estimate $p(x_t=i|\mathbf{y})$, defined as

$$q(x=i|\mathbf{y}, \theta_i) = \theta_i^\top \phi(\mathbf{y})$$

where

$$\theta_i = (\theta_{i,1}, \dots, \theta_{i,B})^\top \in \mathbb{R}^B$$

and $\phi(\mathbf{y})$ is the same as in Eq. (14). The squared loss term in this case is

$$J_i(\theta_i) = \frac{1}{2} \int (q(x=i|\mathbf{y}, \theta_i) - p(x=i|\mathbf{y}))^2 p(\mathbf{y}) d\mathbf{y}.$$

Expanding and using $p(x|\mathbf{y}) = p(\mathbf{y}|x)p(x)/p(\mathbf{y})$ we obtain

$$\begin{aligned} J_i(\theta_i) &= \frac{1}{2} \int q(x=i|\mathbf{y}, \theta_i)^2 p(\mathbf{y}) d\mathbf{y} \\ &\quad - \int q(x=i|\mathbf{y}, \theta_i) p(\mathbf{y}|x=i) p(x=i) d\mathbf{y} + C, \end{aligned}$$

which can be empirically approximated to give the following training criterion:

$$\hat{J}_i(\boldsymbol{\theta}_i) = \frac{1}{2} \boldsymbol{\theta}_i^\top \boldsymbol{\Phi}^\top \boldsymbol{\Phi} \boldsymbol{\theta}_i - \boldsymbol{\theta}_i^\top \boldsymbol{\Phi} \mathbf{m}_i + \frac{\rho}{2} \|\boldsymbol{\theta}_i\|^2.$$

$\hat{J}_i(\boldsymbol{\theta}_i)$ is minimized by

$$\hat{\boldsymbol{\theta}}_i = \left(\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \rho \mathbf{I}_B \right)^{-1} \boldsymbol{\Phi}^\top \mathbf{m}_i, \quad (17)$$

which is essentially kernel ridge regression. As in the two-class case, we select ρ and σ with cross validation. Also as before we round up negative estimates to zero, $q(x=i|\mathbf{y}, \hat{\boldsymbol{\theta}}_i) = \max\left(0, \hat{\boldsymbol{\theta}}_i^\top \boldsymbol{\phi}(\mathbf{y})\right)$, though we note that that negative estimates are unlikely when the number of training samples is large enough [15].

Now using Eq. (16) we are able to calculate likelihood ratios for any pair of states using

$$\hat{w}_{i,j}(t) = \frac{n_i}{n_j} \frac{q(x_t=i|\mathbf{y}_t, \hat{\boldsymbol{\theta}}_i)}{q(x_t=j|\mathbf{y}_t, \hat{\boldsymbol{\theta}}_j)},$$

where n_i is the number of training samples labeled $x_t = i$. Just as in the two-class case, Eq. (17) can be computed very quickly.

4.3 Setting other parameters

We now discuss setting parameters \mathbf{A}, π in the density ratio HMM given training data, first when states $x_{1:T}$ are available and second as an unsupervised learning problem when we only have observation sequences $\mathbf{y}_{1:T}$.

If $x_{1:T}$ is observed in training data, $p(x_t=i|x_{t-1}=j)$ and $p(x_1=i)$ can be estimated directly by frequency. The process of setting the remaining parameters \mathbf{A}, π in the case that the labels $x_{1:T}$ are present in the training data is therefore identical to the standard HMM; we refer the reader to the details in (author?) [12].

For the unsupervised case in which only the observation sequence $\mathbf{y}_{1:T}$ is available for training, learning can be carried out by iterating between (1) a likelihood-maximization step to update $\mathbf{A}, \pi, \{\boldsymbol{\theta}_i\}$ given estimates of $p(x_{1:T}|\mathbf{y}_{1:T})$ and (2) an expectation step, the inference procedure given in §3. To do this some initial estimate $\hat{x}_{1:T}$ is required as a starting point, which could be obtained by running a standard non-dynamic clustering procedure such as k -means on the observations. The iterations are continued until convergence or some limit on the number of cycles is reached. For unsupervised learning we require one alteration to the parameter estimation procedure for $w_{ij}(\cdot)$ in order to accommodate soft estimates of state probabilities γ_i rather than hard labels. The weighted version of Eq. (17) is

$$\hat{\boldsymbol{\theta}}_i = \left(\boldsymbol{\Phi}^\top \boldsymbol{\Gamma}^{(i)} \boldsymbol{\Phi} + \rho \mathbf{I}_B \right)^{-1} \boldsymbol{\Phi}^\top [\gamma_i(1), \dots, \gamma_i(B)]^\top$$

where $\boldsymbol{\Gamma}^{(i)}$ is a $B \times B$ diagonal matrix such that $\boldsymbol{\Gamma}_{j,j}^{(i)} = \gamma_i(j)$. The maximization step updates for \mathbf{A}, π given γ_i are again identical to the standard HMM case [12].

5 Experiments

We now give experimental results using the above methods when applied to both synthetic and real-world datasets.

5.1 Sequential classification on toy data

To give an illustration of HMM inference using density ratio estimation, we generated data from a simple switching model with three dynamical regimes. The transition probabilities in the model were set to

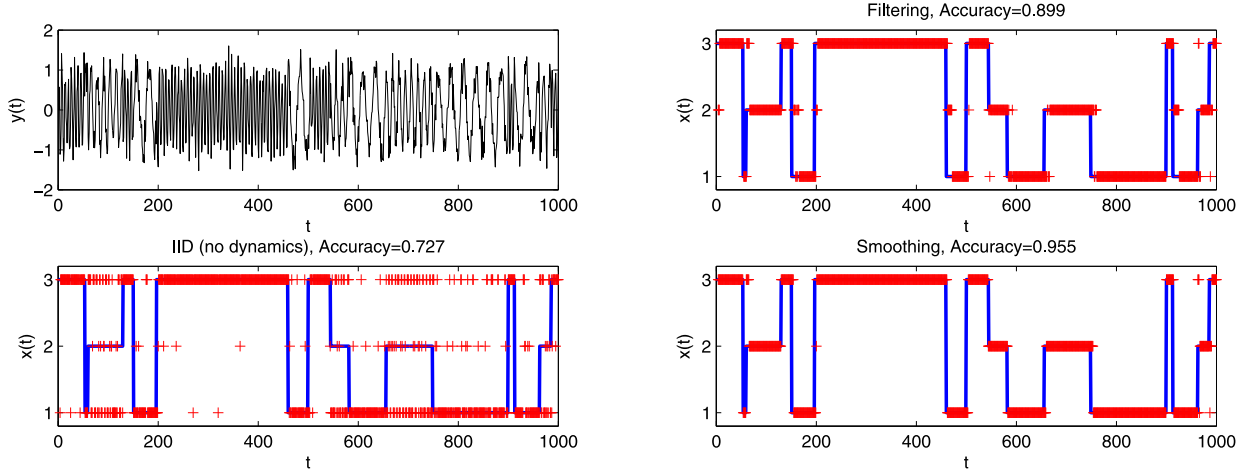


Figure 2: Inference example on noisy sine wave sequence. The top left panel shows a sampled sequence, which switches randomly between three dynamic regimes (sine waves of three different frequencies with Gaussian noise added). In the bottom left panel, the blue line shows the true value of x_t , and the red crosses show the estimated most likely value given density ratio estimates independently at each time frame. The top right panel shows the inference improving when forward dynamics are incorporated; the bottom right panel shows further improvement when both forward and backward dynamics are incorporated.

$\mathbf{A} = \begin{bmatrix} .98 & .01 & .01 \\ .01 & .98 & .01 \\ .01 & .01 & .98 \end{bmatrix}$ with initial state probabilities $\pi = [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]$, and these parameters were then used to generate sequences $x_{1:T}$. Scalar sequences $y_{1:T}$ were then generated conditioned on x_t as follows:

$$y_t = \begin{cases} \sin(0.2t) + \eta_t & x_t = 1 \\ \sin(0.4t) + \eta_t & x_t = 2 \\ \sin(0.6t) + \eta_t & x_t = 3 \end{cases}$$

with $\eta_t \sim \mathcal{N}(0; 0.25)$. A sample sequence is shown in Figure 2 (top left). To construct the vector sequence $\mathbf{y}_{1:T}$ used for testing inference, we took subsequences of $y_{1:T}$ using a sliding window of length d_y (using $d_y = 4$ for this example) such that $\mathbf{y}_t = y_{t-d_y+1:t}$, for $t \geq d_y$.

Figure 2 (bottom left) shows the results of using density ratios at each time step to find the most likely class independently at each time step, by finding $\arg \max_i \left(\sum_{j \in \mathcal{S}} r_{i,j}(\mathbf{y}_t) \right)$. This is equivalent to filtering inference using a uniform transition matrix $\mathbf{A}_{ij} = \frac{1}{S}$, and gives an idea of how informative the subsequences \mathbf{y}_t are about x_t at individual time frames when no dynamical information is incorporated. In this plot the blue line shows the true values of $x_{1:T}$ in the sample sequence, and the red crosses show the MAP estimates. The top right panel shows MAP estimates with filtering inference using the forward equations only. The bottom right panel shows MAP estimates with smoothing inference using both forward and backward steps.

Inference results were compared to those from two alternative models. The first was an alternative nonparametric hidden Markov model, using kernel density estimation (using a squared exponential kernel, with parameters chosen by cross-validation) to model $p(\mathbf{y}_t | x_t = i)$ for each $i \in \mathcal{S}$, as proposed in (author?) [8]. The second was the conventional Gaussian mixture model approach to modeling the observation density of each regime [12], with the number of components of the mixture selected in each case being that which minimized the Bayesian Information Criterion (BIC) on training data. Training sequences were randomly generated from the above switching noisy sine wave example, of lengths between 50 and 500 subsequences for each of the three dynamic regimes. Test sequences of length 1000 were also generated as above. The density ratio HMM, KDE-HMM and GMM-HMM were trained and applied to the test sequence. The accuracy,




Figure 3: Accuracy of state estimates in noisy switching sine wave problem, mean and standard deviation over 500 runs per training sample size.

calculated as the proportion of time frames for which the MAP estimate $\max_i p(x_t=i|\mathbf{y}_{1:T})$ was equal to the true simulated state x_t , was calculated for both methods. This was repeated 500 times for each training sample size. Figure 3 shows the mean and standard deviation of accuracy for filtering and smoothing with all three methods. The density ratio HMM has consistently higher average accuracy than the other methods, particularly for small training set sizes.

5.2 Physiological monitoring

We next evaluated the density ratio HMM using time series of physiological measurements from premature infants receiving intensive care. The dataset we used consists of 24 hour sequences of monitoring data from 15 babies, a total of 360 hours of data with measurements taken at one second intervals. The measurements are of vital signs (heart rate, blood pressures, temperatures, blood gas concentrations) and environmental measurements (incubator temperature and humidity). The data is annotated with the occurrences of four common phenomena: bradycardia (temporary slowing or stopping of the heart), opening of the incubator, the taking of a blood sample, and disconnection of the core temperature probe. Any period in the data during which there was some clinically significant change not covered by one of the four conditions above was annotated as a fifth class. Finally, for each baby a period of around 10 minutes was annotated as ‘normal’, i.e. representative of that baby’s baseline physiology. The data is publicly available and described in (author?) [11].

We first trained a set of two-class density ratio HMMs with this data, treating each state in the annotation as a separate inference problem. To train the bradycardia model for a particular baby, for example, we would take the period for that baby annotated as normal as training data for the first state, and periods annotated as bradycardia from other babies as the training data for the second state.

We compared the output of the density ratio HMM to that of a factorial hidden Markov model (FHMM) and factorial switching linear dynamical system (SLDS), recreating the evaluation of (author?) [11]. As the problem associated with this dataset is real-time patient monitoring, we applied filtering inference only. The latter methods were developed using extensive domain knowledge as to the physical processes underlying the observations. The evaluation was done using 3-fold cross validation on the set of 15 sequences. Area under ROC curve (AUC) and equal error rate (EER) for each of the classes in the annotations were calculated for the three methods, shown in Table 5.2. The density ratio HMM gives either equivalent or superior results in all cases, for example achieving a 17% increase in AUC and 13% decrease in EER for detection of temperature probe disconnection compared to the next best method.

		FHMM	SLDS	DR-HMM
Bradycardia	AUC	.66	.88	.92
	EER	.37	.25	.13
Incu. Open	AUC	.78	.87	.88
	EER	.25	.17	.17
Blood Sample	AUC	.82	.96	.96
	EER	.20	.14	.05
Temp. Probe	AUC	.74	.77	.94
	EER	.32	.23	.10
Abnormal	AUC	-	.69	.75
	EER	-	.36	.31

Table 1: Classification accuracy on neonatal intensive care unit time series data, for occurrences of bradycardia, opening of the incubator, blood sampling, temperature probe disconnection and “other significant deviation from normal dynamics”. FHMM denotes the factorial hidden Markov model, SLDS denotes a switching linear dynamical system, and DR-HMM denotes the density ratio hidden Markov model. The highest AUC and lowest EER are highlighted in bold.

Anomaly detection

We also constructed an anomaly detection model for each test sequence, to assess the ability of this framework to identify any clinically significant deviation from known types of physiological variation. Using the multiple-class density ratio estimation procedure in §4.2, this can be modeled quite easily¹.

Assuming that observations from outlier classes might be present in test data, we use $x_t = *$, $*$ $\notin \mathcal{S}$ to denote any such class at time t . The method we propose for discriminating outliers from inliers is similar in essence to the one-class support vector machine [14] and the kernel Fisher discriminant method for outlier detection [13]. These methods are based on the assumption that outliers occupy low-density regions of the data space and that a kernel model can be used to characterize the high-density regions given training data. Any given significance threshold can then be used to separate the inlier and outlier level sets.

We estimate the conditional probability of an outlier $p(x = * | \mathbf{y}, \boldsymbol{\theta}_i)$ with

$$q(x = * | \mathbf{y}, \boldsymbol{\theta}_*) = 1 - \boldsymbol{\theta}_*^\top \boldsymbol{\phi}(\mathbf{y}). \quad (18)$$

The problem of identifying outliers can then be equated with learning $\boldsymbol{\theta}_*$ such that Eq. (18) is close to zero when \mathbf{y} is within a region in which training data has high density, and is close to one anywhere else. To achieve this we minimize the following loss function:

$$J_2(\boldsymbol{\theta}_2) = \frac{1}{2} \int \left(1 - \boldsymbol{\theta}_2^\top \boldsymbol{\phi}(\mathbf{y})\right)^2 p(\mathbf{y}) d\mathbf{y} + \frac{\rho}{2} \|\boldsymbol{\theta}_2\|^2. \quad (19)$$

The solution to this is given by

$$\begin{aligned} \hat{\boldsymbol{\theta}}_* &= \left(\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \rho \mathbf{I}_B\right)^{-1} \sum_{t=1}^N \boldsymbol{\phi}(\mathbf{y}_t) \\ &= \left(\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \rho \mathbf{I}_B\right)^{-1} \sum_{i \in \mathcal{S}} \boldsymbol{\Phi}^\top \mathbf{m}_i \\ &= \sum_{i \in \mathcal{S}} \hat{\boldsymbol{\theta}}_i. \end{aligned}$$

¹An outlier detection method which could be used with the two-class ratio estimation method in §4.1 is described in (author?) [4].

Therefore

$$q(x = * | \mathbf{y}, \hat{\boldsymbol{\theta}}_1, \dots, \hat{\boldsymbol{\theta}}_S) = 1 - \sum_{i \in \mathcal{S}} \hat{\boldsymbol{\theta}}_i^\top \boldsymbol{\phi}(\mathbf{y}).$$

The parameters learned to model the inlier classes can therefore also be used for outlier detection.

To apply this to the problem of anomaly detection in the physiological monitoring dataset, we trained inlier class parameters using the reference period of the test sequence annotated as normal as well as any subsequence of training data annotated as bradycardia, incubator opening, blood sampling or temperature probe disconnection. Inference of $p(x_{1:T} = * | \mathbf{y}_{1:T})$ in test data therefore gave an estimate of whether any significant physiological changes not consistent with any of the known dynamical regimes were occurring.

We compared this to the anomaly detection method based on the SLDS in (author?) [11]. Again using 3-fold cross-validation on the 360 hours of annotated data, the performance of our method was found to be superior to that of the SLDS, with 6% increase in AUC and 5% decrease in EER.

6 Conclusions

In this paper we have demonstrated that direct density ratio estimation methods can be applied to sequential inference problems, improving the discriminative capability of HMMs without losing the probabilistic interpretation of such models. We believe this method is particularly effective when no parametric class-conditional distribution of observations is known a priori, which is usually the case in real-world problems. As density ratio estimation is a growing field, having already been successfully applied to various problems in statistical inference such as covariate shift and outlier detection, the ideas in this paper make further advances in the field directly applicable to sequence modeling. It would be possible to apply this same principle to several other related sequential probabilistic models, another significant direction for future work.

References

- [1] Y. Altun, I. Tsochantaridis, and T. Hofmann. Hidden Markov support vector machines. In *Proceedings of the 20th International Conference on Machine Learning*, 2003.
- [2] S. Bickel, M. Brückner, and T. Scheffer. Discriminative learning for differing training and test distributions. In *Proceedings of the 24th Annual International Conference on Machine Learning (ICML2007)*, pages 81–88, 2007.
- [3] A. Gretton, A. Smola, J. Huang, M. Schmittfull, K. Borgwardt, and B. Schölkopf. Covariate shift by kernel mean matching. In J. Quiñero-Candela, M. Sugiyama, A. Schwaighofer, and N. Lawrence, editors, *Dataset Shift in Machine Learning*, pages 131–160, Cambridge, MA, USA, 2009. MIT Press.
- [4] T. Kanamori, S. Hido, and M. Sugiyama. A least-squares approach to direct importance estimation. *Journal of Machine Learning Research*, 10:1391–1445, Jul. 2009.
- [5] T. Kanamori, T. Suzuki, and M. Sugiyama. Statistical Analysis of Least-Squares Density Ratio Estimation. *Machine Learning*, 86(3):335–367, 2012.
- [6] A. McCallum, D. Freitag, and F. Pereira. Maximum entropy markov models for information extraction and segmentation. In *Proceedings of the Seventeenth International Conference on Machine Learning*, volume 951, pages 591–598, 2000.
- [7] X. Nguyen, M. J. Wainwright, and M. I. Jordan. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, 2010.
- [8] M. Piccardi and O. Perez. Hidden Markov models with kernel density estimation of emission probabilities and their use in activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2007.

- [9] V. Punyakanok and D. Roth. The Use of Classifiers in Sequential Inference. In *Advances in Neural Information Processing Systems*, volume 14, 2001.
- [10] J. Qin. Inferences for case-control and semiparametric two-sample density ratio models. *Biometrika*, 85(3):619–630, 1998.
- [11] J.A. Quinn, C.K.I. Williams, and N. McIntosh. Factorial switching linear dynamical systems applied to physiological condition monitoring. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31:1537–1551, 2009.
- [12] L.R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [13] V. Roth. Kernel fisher discriminants for outlier detection. *Neural computation*, 18(4):942–960, 2006.
- [14] B. Schölkopf, R. Williamson, A. Smola, and J. Shawe-Taylor. SV estimation of a distribution’s support. In *Advances in Neural Information Processing systems*, volume 12, 1999.
- [15] M. Sugiyama, H. Hachiya, M. Yamada, J. Simm, and H. Nam. Least-squares probabilistic classifier: A computationally efficient alternative to kernel logistic regression. In *Proc. International Workshop on Statistical Machine Learning for Speech Processing (IWSML2012)*, pages 1–10, 2012.
- [16] M. Sugiyama, T. Suzuki, and T. Kanamori. *Density Ratio Estimation in Machine Learning*. Cambridge University Press, Cambridge, UK, 2012.
- [17] M. Sugiyama, T. Suzuki, and T. Kanamori. Density ratio matching under the Bregman divergence: A unified framework of density ratio estimation. *Annals of the Institute of Statistical Mathematics*, 64(5):1009–1044, 2012.
- [18] M. Sugiyama, T. Suzuki, S. Nakajima, H. Kashima, P. von Büna, and M. Kawanabe. Direct importance estimation for covariate shift adaptation. *Annals of the Institute of Statistical Mathematics*, 60(4):699–746, 2008.
- [19] B. Taskar, C. Guestrin, and D. Koller. Max-margin Markov networks. In *Advances in Neural Information Processing Systems*, volume 16, 2004.
- [20] I. Tsochantaridis, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research*, 6(2):1453, 2006.
- [21] V.N. Vapnik. *Statistical Learning Theory*. Wiley, 1998.
- [22] P.C. Woodland and D. Povey. Large scale discriminative training of hidden Markov models for speech recognition. *Computer Speech & Language*, 16(1):25–47, 2002.